

1        SYSTEM AND METHOD FOR SELECTIVELY INCREASING MESSAGE

2                                TRANSACTION COSTS

3  
4                                BACKGROUND OF THE INVENTION

5        1.        Field of the Invention

6                The present invention relates generally to systems and methods for  
7        transporting messages, and more particularly to a system and method for selectively  
8        increasing transaction costs.

9        2.        Discussion of Background Art

10              Mass mailings, also known as junk mail, received through postal mail carriers  
11        is a pervasive part of most peoples lives. A desired benefit to those who originate  
12        such communications is an increase in revenues from the sale of products or services  
13        advertised in the mailing. Such physical mass mailings received by most individuals  
14        is limited, due to the transaction costs associated with their creation and mailing.

15              In contrast, mass e-mailings containing such advertisements are much less  
16        expensive to create and send than physical mail. As a result, the number of electronic  
17        mass mailings received by most individuals is exponentially growing, since the  
18        transaction costs are essentially mainly at the recipient's and not the sender's end.  
19        Many e-mail users receive hundreds of such junk e-mailing mailings a day. Such  
20        advertisements tend to obscure the few e-mails of real importance to a user. A term  
21        coined for such mass e-mailings is "spam." Not only does spam waste the time of  
22        most users, it's increasingly consuming a significant amount of network bandwidth  
23        and thus is shifting some of the transaction costs to service providers as well.

24              There have been some proposals to deal with spam, but they have their  
25        limitations. For instance, user spam e-mail filtering programs have problems such as  
26        false hits, resulting in real messages being treated as spam, and false misses, resulting

1 in a significant amount of spam not getting filtered out. Filtering based on approved  
2 email addresses often prevents a user from receiving desired email from people not on  
3 the list. Suing spammers for consuming network resources has largely failed,  
4 especially for spammers in other countries. Several approaches have been proposed  
5 that make spam costly. One scheme requires paying for each email and having the  
6 money refunded if the mail is accepted, but there is no agreed upon form of money  
7 that can be used. Another scheme only accepts mail if the sender can prove that an  
8 expensive computation was done, but this approach requires changes to a user's e-  
9 mail program as well as the Simple Mail Transfer Protocol (SMTP) for transporting e-  
10 mail.

11 In response to the concerns discussed above, what is needed is a system and  
12 method that overcomes the problems of the prior art.

13



1                                    BRIEF DESCRIPTION OF THE DRAWINGS

2                    Figure 1 is a dataflow diagram of one embodiment of a system for increasing  
3    message transaction costs;

4                    Figure 2 is a flowchart of one embodiment of a first method for increasing  
5    message transaction costs;

6                    Figure 3 is a flowchart of one embodiment of a second method for increasing  
7    message transaction costs; and

8                    Figure 4 is a flowchart of one embodiment of a third method for increasing  
9    message transaction costs.

10

1           DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

2           The present invention modifies the economics of spamming by selectively  
3           increasing the transaction costs of sending e-mail. This is a process of shifting the  
4           transaction costs of sending spam back on to the sender and off of the network service  
5           providers and final recipients. While such increased costs may have a minimal effect  
6           on legitimate e-mail message transport, such costs will have a significant effect on  
7           those senders of mass junk/spam e-mailings (a.k.a. "spammers").

8           The present invention shifts such transactions costs back on to the senders, by  
9           taking advantage of a difference in the volume of and way spam is delivered, as  
10          compared with legitimate e-mails. Some of the higher transactions costs imposed on  
11          spammers include, requiring a spammer to maintain state information about which  
12          connections have not yet been accepted during repeated message retries should the  
13          spammer wish to continue to send spam to an intended recipient. Such state  
14          information retention would require a spammer to increase their computing, storage,  
15          and network resources/costs. Since spammers often do not wish to incur such costs,  
16          the spammer will typically give up trying to spam addresses protected by the present  
17          invention.

18  
19          Figure 1 is a dataflow diagram of one embodiment of a system 100 for  
20          increasing message transaction costs. Figure 2 is a flowchart of one embodiment of a  
21          first method 200 for increasing message transaction costs. Figures 1 and 2 are now  
22          discussed together.

23          To begin, in step 202, a sender 102 generates an e-mail message on a sender's  
24          computer 104 to be sent over a network 106 to one or more recipients 108. A  
25          sender's Mail User Agent (MUA) 110 operating on the sender's computer 104  
26          actually accepts the sender's data and generates an e-mail message. MUAs allow

1 users to access and manage e-mail, including reading, composing, deleting, printing  
2 and displaying e-mail messages. The sender's MUA 110 also organizes the sender's  
3 data into a format compatible with a Simple Message Transfer Protocol (SMTP). The  
4 SMTP protocol almost universally used for transporting e-mail over a worldwide  
5 network.

6 In step 204, a relaying Message Transfer Agent (MTA) 112 operating on a  
7 network server 114 receives the message and an intended set of recipient addresses  
8 over a data link. For the purposes of this specification, the phrase "a data link" is  
9 herein defined as any portion of a communication path between the sender 102 and the  
10 recipient 108 over which messages are transmitted. For example, the data link, in a  
11 narrow use, can refer to a communication path between a particular sender and a  
12 particular recipient. Alternatively, the data link, a more broad use, can refer to a  
13 communication path between a particular relaying MUA and a particular receiving  
14 MUA.

15 In step 206, the relaying MTA 112 attempts to locate and contact a receiving  
16 MTA 116 over the network 106 which can deliver the message to one or more  
17 recipients in the set of recipient addresses. Communications over the network 106  
18 between the relaying MTA 112 and the receiving MTA 116 conform not only to the  
19 SMTP protocol, but also to TCP/IP and ISO/OSI protocols as necessary.

20 More specifically, MTAs operate as application programs on servers  
21 throughout the network 106. MTA applications map to an "application layer" (layer  
22 7) of the ISO/OSI model. MTA applications incorporate SMTP code for transmitting  
23 e-mail messages. The SMTP code interfaces with the network according to a TCP/IP  
24 protocol. In the TCP/IP protocol. The Transmission Control Protocol (TCP) maps to  
25 a "transport layer" (layer 4) of the ISO/OSI model, and the Internet Protocol (IP) maps  
26 to a "network layer" (layer 3) of the ISO/OSI model. Note that the ISO/OSI model

1 does not map exactly to network communications following the TCP/IP protocol, so  
2 layers 5 and 6 are usually left out. As is discussed below, important features of the  
3 present invention operate at different layers within the ISO/OSI model, and thus  
4 provide many different opportunities for blocking spam.

5 Note that while this discussion bifurcates the sender's computer 104 from the  
6 relaying MTA 112 on the network server 114, the sender's computer 104 and the  
7 server 114 could also be generally labeled as a sending computer. Also, note that  
8 while the recipient's computer 128 is bifurcated from the receiving MTA 116 on the  
9 firewall server 118, and the local server 122, these items 118, 122, and 128 could also  
10 be generally labeled as a receiving computer. Thus, those skilled in the art will  
11 recognize that the exact term used and/or which computer/server contains the  
12 functionality for effecting the present invention can be varied with each embodiment.

13 In step 208, the receiving MTA 116, which in the embodiment herein  
14 described happens to be within a firewall server 118, initializes a database 120 data  
15 structure containing the following fields: contact IP address; a first contact time; and a  
16 randomized delay period.

17 In step 210, upon a first contact within a predetermined time period over the  
18 data link from the relaying MTA 112, the receiving MTA 116: first, populates the  
19 contact IP address field with the IP address of the relaying MTA 112; second,  
20 populates the first contact time field with a time when the relaying MTA 112 first  
21 contacted the receiving MTA 116; and third, populates the randomized delay period.  
22 The first contact time marks the start of the predetermined time period, which is  
23 preferably one day, but which could also be of any length.

24 In step 212, the receiving MTA 116 then transmits a TCP layer "FIN"  
25 command back over the data link to the relaying MTA 112. In step 214, receipt of the  
26 FIN command by the relaying MTA 112 forces the relaying MTA 112 to



1 acknowledging that the data link with the receiving MTA 116 no longer exists. Since  
2 the FIN command is sent at the TCP layer, relaying MTAs cannot ignore such a  
3 command. In contrast, “error codes” transmitted at the SMTP layer in some cases can  
4 be ignored by relaying MTA’s and/or could cause confusion at a relaying MTA  
5 sending legitimate (i.e. non-spam) e-mail messages if the legitimate relaying MTA is  
6 not equipped to handle error codes and/or is not fully SMTP compliant.

7 In step 216, in response to receipt of the “FIN” command, the relaying MTA  
8 112 must either re-queue all messages to be sent over the data link to recipients  
9 through the receiving MTA 116 for later delivery or must give up and delete the e-  
10 mail messages to be sent. Re-queuing is a standard step for relaying MTAs sending  
11 legitimate e-mails, however, re-queuing would be very resource intensive for a  
12 relaying MTA sending a great number of spam e-mails. As a result, such spamming  
13 relaying MTAs will tend to delete the spam e-mails, thus effectively blocking the  
14 spam from reaching the intended recipients.

15 In step 218, the receiving MTA 116 refuses subsequent connection attempts  
16 over the data link by the relaying MTA 112, if the randomized delay period after the  
17 first contact time has not yet expired.

18 In step 220, once the randomized delay period has expired, the receiving MTA  
19 116 maintains the data link and thus accepts e-mail messages from the relaying MTA  
20 112 for the remainder of the predetermined time period. The receiving MTA 116 then  
21 transfers the messages to a recipient’s MUA 126 within the recipient’s computer 128.  
22 The recipient 108 can then read the message.

23 In step 222, the receiving MTA 116 deletes the relaying MTA’s 116  
24 information in the database 120 after the randomized delay period has expired. The  
25 deletion of the relaying MTA’s 116 information from the database allows for a new  
26 delay period to be calculated once the predetermined time period has expired. While



1 as mentioned above, the predetermined time period is one day, other embodiments of  
2 the present invention may recalculate the delay period either more or less often than  
3 once per day. Preferably, however, the delay period should be calculated at least  
4 every 12 hours in view of a maximum resend threshold of 12 hours observed on the  
5 Internet. Also, the interval in which the delay period is recalculated should also not be  
6 reduced to a point where normal e-mail conversations are excessively interfered with.  
7 Since, spam is sent very broadly, whereas normal e-mail conversations tend to be  
8 much more narrowly addressed and exchanged, spammers will be more affected by a  
9 delay period than regular e-mail senders.

10

11 Figure 3 is a flowchart of one embodiment of a second method 300 for  
12 increasing message transaction costs. Figures 1 and 3 are now discussed together.  
13 The second method 300 is effective against senders of spam messages since the  
14 method 300 takes advantage of how spammers create recipient lists. More  
15 specifically, e-mail list creation programs used by spammers tend to cast a broad net  
16 in order to obtain as many e-mail addresses as possible. Interspersed within such  
17 addresses are many addresses which are either unknown or no longer valid. Currently,  
18 when relaying MTAs send message destined for an unknown/invalid addresses,  
19 receiving MTAs typically reply with a SMTP protocol error 550 "Recipient RCPT  
20 Unknown." The receiving MTAs however continue to permit the relaying MTAs to  
21 other e-mail addresses. The present invention, however, does not permit the relaying  
22 MTA 112 to continue sending messages to the receiving MTA 116, but instead  
23 penalizes the relaying MTA 112 for trying to send an e-mail message to either an  
24 unknown or invalid address.

25 More specifically, in step 302, upon receiving a request to send an e-mail  
26 message to either an unknown or invalid address, the receiving MTA 116 first checks

1 locally for an e-mail address specified by the RCPT command sent by the relaying  
2 MTA 112. In step 304, the receiving MTA 116 then checks with any other local  
3 servers, such as local server 122, residing behind the firewall server 118, by sending  
4 the specified e-mail address along with an "EXPN" command to a final MTA 124  
5 within the local server 122. The EXPN is an SMTP "expand" command requesting  
6 additional information associated with the specified e-mail address, such as the  
7 addressee's full name.

8 Note that the EXPN command is only one of many different methods for  
9 determining the validity of an email address. Those skilled in the art are aware of  
10 other mechanisms that may function equally well, such as the use of RCPT protocol  
11 commands within SMTP, or site specific configurations such as IDENT, or LDAP to  
12 identify valid users.

13 In step 306, if the specified e-mail address is neither found locally on the  
14 receiving MTA 116 or on the local server 122 by the final MTA 124, the receiving  
15 MTA 116 silently hangs-up the network connection with the relaying MTA 112. The  
16 silent hang-up is preferably done at the IP layer within the TCP/IP protocol by adding  
17 a rule to the receiving MTA's 116 firewall code which closes the TCP file handle  
18 without generating any "error codes" or sending any sort of message to the relaying  
19 MTA 112.

20 In step 308, the relaying MTA 112 upon detection that the network connection  
21 is no longer active, automatically generates a TCP timeout. TCP timeouts are hard-  
22 coded into most TCP/IP network chips and, in step 310 automatically prevent the  
23 relaying MTA 112 from attempting to reconnect with the receiving MTA 116 until a  
24 TCP timeout period, typically five minutes, has expired.

25 In contrast to the first method 200 discussed above, the TCP timeout is  
26 generated at the relaying MTA 112 every time the receiving MTA 116 silently drops

1 the data link. Preferably, the receiving MTA 116 silently drops the data link every  
2 time a request is made to transmit an e-mail to an unknown or invalid address, and not  
3 just once per day, like the first method 200. Legitimate e-mail relaying MTAs will in  
4 most cases not be affected by the second 300 method since almost all of the e-mail  
5 addresses relayed by legitimate e-mail relaying MTAs will be valid. However, spam  
6 e-mail recipient lists, contain a disproportionately large number of unknown and  
7 invalid e-mail addresses, and thus will be much more affected and hindered by the  
8 second method 300. Thus spammers will tend to experience many annoying TCP  
9 timeouts each day.

10

11 Figure 4 is a flowchart of one embodiment of a third method 400 for  
12 increasing message transaction costs. Figures 1 and 4 are now discussed together.  
13 The third method 400 is effective against senders of spam messages since the method  
14 400 enhances the second method's 300 effect on spammers. More specifically, since  
15 as mentioned with respect to the second method 300, the e-mail list creation programs  
16 used by spammers tend to cast a broad net in search of e-mail addresses to spam, the  
17 third method 400 creates a sort of "Trojan horse" set of faux e-mail addresses as  
18 spammer bait. These faux addresses can either be treated immediately as invalid, thus  
19 invoking a TCP timeout, or can be treated as "valid" for a predetermined time, after  
20 they are treated as the invalid addresses. Also, different sets of faux addresses can be  
21 created from time to time for confounding senders of spam e-mail.

22 More specifically, in step 402, the receiving MTA 116 within the firewall  
23 server 118, generates a set of faux e-mail addresses. In step 404, the receiving MTA  
24 116 makes the faux addresses publicly available. One way of making the faux  
25 addresses available is to publish them on the network 106. In step 406, the relaying  
26 MTA 112 includes the faux addresses within a spam e-mail recipient list. In step 408,

1 the receiving MTA 116 receives an e-mail message addressed to one or more of the  
2 faux e-mail addresses.

3 Since some spam senders attempt to validate an e-mail address before either  
4 sending spam e-mail to the address or selling the address to other spammers, the  
5 receiving MTA 116, in step 410, provides a faux validation of the faux address back  
6 to the relaying MTA 112.

7 One technique spammers may use to determining if an e-mail address is  
8 “valid” is to insert a tag reference into an HTML portion of the e-mail message. Such  
9 a tag could be either an image reference, a document reference, or some other sort of  
10 reference which requests that information be downloaded from the network 106 or the  
11 relaying MTA 112. In this way a spam sender can indirectly detect that the e-mail  
12 message was opened due to a file downloaded with was pointed to by the tag  
13 reference. Since such tags can be made invisibly small, the recipient 108 may not  
14 even be aware that their opening of the spam e-mail was remotely detected by the  
15 sender 102. Thus, in response to this validation technique, the receiving MTA 116, in  
16 step 412, uses an HTML interpreter to look for and downloads files pointed to by any  
17 tag references in the e-mail message. This gives the relaying MTA 112 an impression  
18 that the e-mail address is valid and was opened by a real user. In step 414, the  
19 receiving MTA 116 then deletes the e-mail and any of the downloaded files.

20

21 After receiving an e-mail directed to one of the faux addresses, the receiving  
22 MTA 116 can respond using one or more of the techniques described below:

23 In a first response technique, in step 416, the receiving MTA 116 immediately  
24 treats the faux addresses as invalid and silently hangs-up the network connection as  
25 described with respect to the second method 300, thereby provoking a TCP timeout in  
26 the relaying MTA 112.

1           In a second response technique, in step 418, the receiving MTA 116 treats the  
2   faux addresses as valid for a predetermined period of time, afterwhich the receiving  
3   MTA 116 treats the faux addresses as invalid, as described with respect to the second  
4   method 300. This technique encourages the relaying MTA 112 to assign a high-  
5   validity label to the faux addresses and thus perhaps permit the faux addresses to  
6   infiltrate a larger number of spam recipient address lists, perhaps including one or  
7   more different spam senders. Spammers often sell their address lists to other  
8   spammers.

9           In a third response technique, in step 420, the receiving MTA 116 treats each  
10   of the faux addresses as valid until a predetermined number of e-mail messages are  
11   sent to that faux address, afterwhich the receiving MTA 116 treats that faux address as  
12   invalid, as described with respect to the second method 300. In step 422, the  
13   receiving MTA 116 then waits until a number of e-mail messages sent to that faux  
14   address drops below a predetermined level, afterwhich the receiving MTA 116 treats  
15   that faux address as valid again. This third technique has an effect similar to the  
16   technique described in step 418, except that a point at which a first spammer sells or  
17   gives the faux address to other spammers is more clear due to a dramatic increase in  
18   spam e-mail received by the faux address. The message rate tends to decrease as  
19   described in step 422 because the spammers may slowly weed out the faux address  
20   from their recipient lists in response to the TCP timeouts.

21           In a fourth response technique, in step 424, the receiving MTA 116 generates  
22   new sets of faux e-mail addresses from time to time, which can then be invalidated  
23   using one of the techniques discussed in steps 416 through 422.

24

25           Those skilled in the art recognize that the techniques just discussed may be  
26   incorporated into various embodiments to varying degrees. And, while one or more

1   embodiments of the present invention have been described, those skilled in the art will  
2   recognize that various modifications may be made. Variations upon and  
3   modifications to these embodiments are provided by the present invention, which is  
4   limited only by the following claims.